

Level 3: Advanced

All Librarians

Module 14

Automating repetitive tasks

The first library automation course designed for nonprogrammers. You'll learn to identify what's worth automating, build your first real automations using visual tools, and see concrete time savings in your actual workflow.

By Yulia Brusova

15 min read · Reviewed June 2026

AI for Academic Libraries · ai-in-academic-libraries.vercel.app

© 2026 · Licensed under CC BY-NC-SA 4.0

WHAT YOU'LL BE ABLE TO DO

- 1 Identify three to five tasks in your current workflow that are worth automating
- 2 Build a basic automation connecting two library tools using Zapier or Make
- 3 Apply AI to a batch processing task that previously required manual effort
- 4 Calculate the time ROI of an automation you've built
- 5 Understand when automation is appropriate and when human judgment is required

The most time-consuming library tasks tend to be recurring ones: weekly usage statistics pulled and pasted into a spreadsheet, semester subject guide footers manually updated with new library hours, patron questions copied from a chat system into a tracking document week after week. None of these tasks require a librarian's professional judgment. They require a librarian's time. Automation returns that time to higher-value work, and the tools that make such automation accessible to library professionals without programming backgrounds have matured considerably in the past several years.

How to decide what is worth automating

Not every repetitive task is worth automating, and applying automation indiscriminately creates its own category of problems. There is no doubt that a disciplined decision framework, applied before any tool is opened, prevents wasted setup time and produces automations that are reliable over time. The framework has four components, each of which must be evaluated honestly before proceeding.

Frequency is the first consideration. For example, a task that recurs daily or several times per week justifies the time required to build and test an automation, whereas a monthly task rarely does. Such a calculation is straightforward but frequently overlooked: library professionals often spend two hours building an automation for a task they perform four times a year, which produces no meaningful time savings over any reasonable period.

Consistency is the second consideration, and it is closely related to the question of judgment. Automation handles consistent, predictable processes well. For example, if a task follows the same numbered steps every time, with no judgment calls between them, it is an excellent candidate for automation. Such a task translates directly into a workflow tool's trigger-action structure without requiring the practitioner to account for exceptions. In contrast, tasks that

require evaluating context at each step, deciding what to do based on information that changes from one instance to the next, are not suitable for full automation. They may benefit from AI assistance at specific points, but a human must remain accountable at each decision.

Volume is the third consideration. For example, processing a batch of two hundred catalog records is worth automating; processing five records is not. Indeed, the efficiency gains of automation depend substantially on scale, and a one-time task, however tedious, rarely justifies the setup investment.

Error cost is the fourth and most consequential consideration. In order to evaluate error cost accurately, the practitioner must ask what happens when the automation makes a mistake, because in any automated system operating over time, mistakes will occur. For example, an automation that updates an internal tracking spreadsheet has low error cost, because a human will review the output and can correct any errors before they have further consequences. An automation that sends patron-facing communications has significantly higher error cost, because errors reach patrons before they can be reviewed. Such patron-facing automations require human review steps built deliberately into the workflow design.

A complementary frame from the AI Fluency course identifies three prerequisites for any delegation decision – automation included – that the four-factor framework above assumes but does not name explicitly (Dakan & Feller, Anthropic Academy, 2025). The first is problem awareness: before opening Zapier, Make, or any other tool, the practitioner must be able to state precisely what success looks like. A librarian who cannot describe what the completed automation should produce is not ready to build it; the tool will fill that gap with assumptions that may not match the actual need. The second is platform awareness: understanding which automation tool handles which type of workflow. Make's conditional logic is not a feature Zapier lacks by accident; it reflects a different architecture suited to a different class of problem. Selecting the wrong tool for a task often produces an automation that technically runs but requires more maintenance than it saves. The third is task delegation: identifying which steps of a workflow should be handled by the automation and which must remain with the practitioner. Such a division is rarely obvious in advance. It requires knowing the workflow well enough to identify exactly where judgment enters the process and exactly where it does not.

Zapier for library workflows: starting simply

Zapier is the most accessible automation tool for library professionals without programming backgrounds. It connects applications using a trigger-action structure: when a specified event occurs in one application, a specified action is executed in another. For example, when

a patron submits a research consultation request through a library booking form, Zapier can automatically log that submission to a shared tracking spreadsheet and send a notification to the reference team. Such a workflow replaces manual copy-pasting and ensures that requests are not missed when staff members are out of the office.

In order to build a first automation, the practitioner selects a trigger, which is an event in one application that initiates the workflow, and an action, which is what happens in a second application as a result. Zapier's interface guides the practitioner through this selection with no code required. Additionally, Zapier maintains a library of pre-built integrations with most applications commonly used in library settings, including Google Workspace, Microsoft 365, [LibCal](#), and standard email providers, which means that many library workflows can be automated without any custom configuration.

Several workflow patterns are particularly well suited to Zapier's structure in library contexts. For example, the form-to-spreadsheet-to-notification pattern handles research consultation intake automatically: a patron submits a LibCal or Google Forms request, Zapier logs the submission to a shared spreadsheet, and sends a [Slack](#) message or email to the relevant staff member. Such a pattern typically requires fifteen minutes to configure and no ongoing maintenance once it is running. The email-to-task pattern serves a different operational need: when email arrives in a specific library inbox containing specified keywords, Zapier creates a corresponding task in a project management tool such as [Asana](#), [Trello](#), or [Notion](#). Furthermore, a new-resource-to-announcement pattern allows a library to queue social media posts or newsletter items automatically when new resources are added to a tracking list, provided the data source supports a trigger event.

It is evident that the free Zapier tier, which supports five active Zaps and a limited number of monthly task executions, is sufficient for most light library use cases. The recommended approach is to start with one automation, test it for several weeks, and evaluate its reliability before building additional workflows. Such incremental implementation allows the practitioner to develop confidence with the tool before building anything that carries higher error costs.

Make for more complex library workflows

Make, formerly known as Integromat, handles more complex automation scenarios than Zapier. Its visual interface represents a workflow as a flowchart, which makes it easier to understand what is happening at each step and to identify where errors are occurring when a workflow does not behave as expected. In order to use Make effectively, the practitioner should map out the full workflow before opening the tool, because the visual design process

is considerably faster when the steps are already established on paper or in a simple diagram.

Make is the appropriate choice when a workflow has more than two steps, when it requires conditional logic, or when it involves processing batches of data rather than responding to individual events. For example, an automated new database trial notification workflow might begin with a trigger when a new row is added to a Google Sheet tracking active trials, pass through a filter that checks whether the trial status column reads "Active," use an AI integration to draft an announcement in the library's established voice based on the database name and description columns, and deposit that draft into a Google Doc queue for staff review before any patron-facing communication is sent. Such a workflow reduces a thirty-minute task to a five-minute review, and it does so without removing human judgment from the final approval step.

Conditional logic is Make's most important practical differentiator from Zapier. For example, a single automation might route a patron request differently based on the patron's institutional role: faculty requests proceed to one staff member, student requests to another, and requests from specific departments are flagged for a subject librarian. Such routing logic requires branching, which Make supports through its filter and router modules. Additionally, Make handles data transformation, which allows the practitioner to extract specific fields from a spreadsheet, reformat dates, or combine text elements before passing data to the next step in the workflow.

Make's free tier allows one thousand operations per month. Indeed, the majority of library automations built in Make for internal workflow purposes will not approach this limit under ordinary operating conditions, which makes it a viable starting point before any paid plan is considered.

Batch processing with AI: high-value library applications

Some of the highest-value library automation does not require Zapier or Make at all. It requires a systematic approach to prompting AI with batches of content that would otherwise require manual processing one item at a time. For example, a library that receives ten patron reference questions of similar types each week can develop a response template for each question type and use AI to customize each template for the specific patron's wording and context. Such an approach reduces twenty minutes of response drafting to approximately two minutes of prompt submission and review.

Batch metadata work is another application with a significant time return. In order to normalize inconsistent catalog records, the practitioner exports a CSV of the affected records and pastes batches of thirty to fifty records into [Claude](#) or [ChatGPT](#) with a clear instruction: for each record in this list, suggest a normalized publisher name, and flag any records that appear to be duplicates. Such a prompt produces a structured output that can be reviewed in a fraction of the time required to check each record manually. Furthermore, the AI output is not applied directly to the catalog; it is reviewed and corrected by a staff member before any changes are made, which keeps a human accountable for the final state of the records.

Document summarization is a third high-value batch application. For example, monthly vendor reports, annual assessment survey results, and faculty feedback documents all follow consistent internal structures, which makes them well suited to AI summarization. A prompt requesting a two-paragraph executive summary and a list of action items from a ten-page vendor report typically produces an accurate and usable output in under thirty seconds. Such a workflow is particularly valuable for library professionals who receive reporting documents from multiple vendors or departments and need to synthesize them before a budget or planning meeting.

The time mathematics of batch processing are concrete and worth calculating explicitly. For example, a monthly chat reference statistics summary that requires forty-five minutes when completed manually typically takes under ten minutes when completed with an AI-assisted workflow: the statistics are submitted to Claude, a summary paragraph is requested along with a comparison to the previous month, and the output is reviewed and lightly edited before use. Such a savings of approximately thirty-five minutes per month compounds across all recurring documentation tasks to produce a meaningful return across a full academic year.

When automation is inappropriate and what it cannot replace

Automation is powerful when tasks are consistent, low-stakes, and well-defined. It creates serious problems when applied to tasks that do not meet those conditions. There is no doubt that the professional judgment required to identify which tasks are appropriate for automation is itself a core competency, one that no automation tool can supply or substitute.

Certain categories of library work must not be automated. For example, communications that require empathy or individual judgment, such as responding to a patron in distress or addressing a sensitive reference question, require human presence at every step of the interaction. Such situations call for the practitioner's professional assessment of what the patron needs, and that assessment cannot be reduced to a trigger-action rule without meaningful risk of harm. Similarly, decisions with institutional consequences, including

collection weeding decisions and access policy changes, require human accountability and must not be delegated to an automated system operating without oversight.

Patron-facing communications warrant particular caution regardless of how well an automation has performed in testing. For example, an automation that sends emails directly to patrons without a human review step removes the library's ability to catch errors before they reach the patron. Such errors range from minor, a date formatted incorrectly, to significant, a message sent to the wrong patron or containing inaccurate policy information. In order to use automation in patron-facing contexts responsibly, a human review step must be built into the workflow design as a structural requirement rather than added as an afterthought when problems arise.

It is evident that the most important limitation of automation is not technical but conceptual. Automation handles the standard case. The practitioner handles everything else. For example, an intake form automation can route a patron's consultation request to the correct staff member under normal circumstances, but when the patron's need does not fit a standard category, when the request is unusual or urgent or requires an exception to established procedure, a human must assess and respond. Such exceptions are not failures of the automation; they are exactly the situations that distinguish professional library work from rule-following. Professional judgment is not what remains after automation is applied to routine tasks. It is the core of library practice, and automation is one of several tools that library professionals use in order to protect and preserve time for that work.

Dakan and Feller describe the professional responsibility that applies at the point of sharing AI-assisted work as deployment diligence - the practitioner, not the tool, is accountable for what is produced and distributed (Dakan & Feller, Anthropic Academy, 2025). Such accountability extends without qualification to automated workflows. When an automation generates a patron-facing email, a statistical summary sent to administrators, or a catalog record deposited without manual review, the standard of accuracy expected is identical to the standard that would apply to work completed entirely by hand. The automation is not an explanation for an error; it is the means by which the practitioner chose to produce the work. This framing makes the case for human review steps considerably stronger than a list of risks: review steps exist not because automations are unreliable, but because the practitioner remains professionally responsible for every output that leaves the library under the library's name.

PRACTITIONER NOTE

The automation that changes library workflows most reliably is rarely the most sophisticated one. A simple Zap connecting a LibCal appointment form to a shared tracking spreadsheet, with a Slack notification to the team, takes approximately five minutes to build and prevents missed appointments when staff members are out of the office. The decision about which tasks to automate matters considerably more than the technical complexity of the automations that result.

KEY TAKEAWAYS

- A task worth automating is frequent, consistent in its steps, high in volume, and low in error cost if the automation produces a mistake.
- Zapier connects two applications through a trigger-action structure and is the most accessible starting point for library professionals without programming backgrounds.
- Make handles workflows requiring conditional logic, multiple sequential steps, and batch data processing, and is the appropriate tool when a Zapier Zap is not sufficient.
- Batch AI processing, submitting groups of similar documents, records, or patron questions to an AI tool with a structured prompt, produces significant time savings without requiring any automation platform.
- Patron-facing communications and decisions with meaningful institutional consequences must not be fully automated; human review steps are non-negotiable at those points.
- Professional judgment is not what remains after automation handles routine tasks; it is the core of library practice, and automation exists in order to protect the time available for that work.

References

APA 7TH EDITION

1. Dakan, R., & Feller, A. (2025). *AI fluency: Framework & foundations* [Online course]. Anthropic Academy. <https://anthropic.skilljar.com/ai-fluency-framework-foundations>

ACRL AI Competencies covered

Use & Application

Sub-competencies: 4.1, 4.4 · ACRL AI Competencies (2025)